



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/010,572	11/09/2001	Carl Cavanagh	5181-97900	4344

7590 01/03/2006
Lawrence J. Merkel
Conley, Rose & Tayon, P.C.
P.O. Box 398
Austin, TX 78767

EXAMINER

SHARON, AYAL I

ART UNIT PAPER NUMBER

2123

DATE MAILED: 01/03/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

DETAILED ACTION

Introduction

1. Claims 1-21, 23-30, and 32-33 of U.S. Application 10/010,572, originally filed on 11/09/2001, are currently pending.
2. Claims 15-23 and 30 have been amended.
3. Claims 22 and 31 have been cancelled.
4. Applicants' arguments (see p.10 of Applicants' amendment filed on 10/24/2005), with respect to the rejections of claims 15-23 under 35 USC § 101 have been fully considered and are persuasive in light of the amendments to the claims. Therefore, the rejections have been withdrawn.
5. Applicant's arguments, (see pp.9-14 of Applicants' amendment filed on 10/24/2005), with respect to the rejections of claims 1-21, 23-30, and 32-33 under 35 U.S.C. 102(b) as anticipated by Agrawal_1 have been fully considered and are persuasive. Therefore, the rejection has been withdrawn.
6. However, upon further consideration, new grounds of rejection of claims 1-21, 23-30, and 32-33 are made in view of the Agrawal_1 reference described below.
7. This action is non-final.

Claim Rejections - 35 USC § 102

8. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

9. The prior art used for these rejections is as follows:

- Agrawal, P. et al. "A Hardware Logic Simulation System." IEEE Transactions on CAD of Integrated Circuits and Systems. Jan. 1990. Vol. 9, Issue 1, pp.19-29. (Henceforth referred to as "**Agrawal_1**").

10. The claim rejections are hereby summarized for Applicant's convenience. The detailed rejections follow.

- 11. Claims 1-4, 15-20, 24-29, and 33 are rejected under 35 U.S.C. 102(b) as being anticipated by Agrawal_1.**

12. In regards to Claim 1, Agrawal_1 teaches the following limitations:

1. A distributed simulation system comprising a plurality of nodes,

See Agrawal_1, p.19, right column, 3rd para., which teaches that "The MARS hardware consists of a cluster of 15 processing elements (PE's) connected by a crossbar switch ..."

wherein each node of the plurality of nodes is configured to simulate a different portion of a system under test using a simulator program configured to perform a simulation as a series of timesteps, and

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm", which teaches that "... All other gates or functional boxes in the circuit are not processed."

wherein each timestep includes at least a first phase and a second phase, and

Art Unit: 2123

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm", which teaches a *fan-out update* phase and an *evolution* phase.

wherein each node of the plurality of nodes is configured not to cause the simulator program to evaluate a model of the different portion of the system under test during the first phase even if one or more commands are received by that node during the first phase, and

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm", which teaches that "Processing events in two phases permits all input changes at one gate to be accumulated before the gate is evaluated."

wherein each node of the plurality of nodes is configured to cause the simulator program to evaluate the model during the second phase in response to receiving a command including one or more signal values for signals of the model.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm", which teaches that "The basic philosophy of this technique is to *evaluate* (compute the output as a function of the inputs) a gate or functional unit if and only if an *event* (change in the value of an input signal) occurs at the gate inputs."

13. In regards to Claim 2, Agrawal_1 teaches the following limitations:

2. The distributed simulation system as recited in claim 1 wherein each node of the plurality of nodes is configured not to cause the simulator program to evaluate the model during the second phase if the signal values in the command received by that node are the same as the current values of the signals.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm"

14. In regards to Claim 3, Agrawal_1 teaches the following limitations:

3. The distributed simulation system as recited in claim 1 wherein each node of the plurality of nodes is configured, if one or more output signals of the model change in response to evaluating the model, to transmit a command including at least the signal values of the output signals that change.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm"

15. In regards to Claim 4, Agrawal_1 teaches the following limitations:

4. The distributed simulation system as recited in claim 1 wherein each node of the plurality of nodes is configured to cause the simulator program to evaluate the model two or more times during the second phase in response to two or more commands including signal values.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm"

16. In regards to Claim 15, Agrawal_1 teaches the following limitations:

15. A computer readable medium storing instructions which, when executed on a computer, process a first one or more commands received during a first phase of a timestep without causing a simulator program to evaluate a model, and cause the simulator program to evaluate the model during a second phase of the timestep in response to receiving a second command including one or more signal values for signals of the model.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm", which teaches a *fan-out update* phase and an *evolution* phase.

Section 2.1 of Agrawal_1 also teaches that "Processing events in two phases permits all input changes at one gate to be accumulated before the gate is evaluated", and

Section 2.1 of Agrawal_1 also teaches that "The basic philosophy of this technique is to *evaluate* (compute the output as a function of the inputs) a gate or functional unit if and only if an *event* (change in the value of an input signal) occurs at the gate inputs."

17. In regards to Claim 16, Agrawal_1 teaches the following limitations:

16. The computer readable medium as recited in claim 15 wherein the instructions, when executed, do not cause the simulator program to evaluate the model during the second phase if the signal values in the second command are the same as the current values of the signals in the model.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm"

18. In regards to Claim 17, Agrawal_1 teaches the following limitations:

17. The computer readable medium as recited in claim 15 wherein the instructions, when executed, if one or more output signals of the model change in response to evaluating the model, transmit a command including at least the signal values of the output signals that change.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm"

19. In regards to Claim 18, Agrawal_1 teaches the following limitations:

18. The carrier medium as recited in claim 17 wherein the instructions, when executed, if no output signals change value during the second phase, transmit a no-operation command.

Agrawal_1 teaches sending the "zero message" at the end of the evaluation phase "down the pipeline to reconfigure the pipeline for the next fan-out update phase". (See Agrawal_1, p.21, "2.2.2 Evaluation Phase", 1st para.).

20. In regards to Claim 19, Agrawal_1 teaches the following limitations:

19. The computer readable medium as recited in claim 15 wherein the instructions, when executed, cause the simulator program to evaluate the model two or more times during the second phase in response to two or more commands including signal values and optional signal strengths.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm"

21. In regards to Claim 20, Agrawal_1 teaches the following limitations:

20. The computer readable medium as recited in claim 15 wherein, in response to a third command indicating an end of the first or second phase, is configured to return to the simulator program.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm"

22. In regards to Claim 24, Agrawal_1 teaches the following limitations:

24. A method comprising:

receiving a first one or more commands in a node of a distributed simulation system during a first phase of a timestep;

processing the first one or more commands without causing a simulator program to evaluate a model;

receiving a second command during a second phase of the timestep; and

processing the second command including causing the simulator program to evaluate the model if the second command includes one or more signal values for signals of the model.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm", which teaches a *fan-out update* phase and an *evolution* phase.

Section 2.1 of Agrawal_1 also teaches that "Processing events in two phases permits all input changes at one gate to be accumulated before the gate is evaluated", and

Section 2.1 of Agrawal_1 also teaches that "The basic philosophy of this technique is to *evaluate* (compute the output as a function of the inputs) a gate or functional unit if and only if an *event* (change in the value of an input signal) occurs at the gate inputs."

23. In regards to Claim 25, Agrawal_1 teaches the following limitations:

25. The method as recited in claim 24 wherein processing the second command does not include causing the simulator program to evaluate the model if the signal values in the second command are the same as the current values of the signals in the model.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm"

24. In regards to Claim 26, Agrawal_1 teaches the following limitations:

26. The method as recited in claim 24 further comprising, if the evaluation of the model during the second phase results in one or more output signals of the model changing, transmitting a command including at least the signal values of the output signals that change.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm"

25. In regards to Claim 27, Agrawal_1 teaches the following limitations:

27. The method as recited in claim 26 further comprising, if no output signals change value during the second phase, transmitting a no-operation command.

Agrawal_1 teaches sending the "zero message" at the end of the evaluation phase "down the pipeline to reconfigure the pipeline for the next fan-out update phase". (See Agrawal_1, p.21, "2.2.2 Evaluation Phase", 1st para.).

26. In regards to Claim 28, Agrawal_1 teaches the following limitations:

28. The method as recited in claim 24 further comprising causing the simulator program to evaluate the model two or more times during the second phase in response to two or more commands including signal values.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm"

27. In regards to Claim 29, Agrawal_1 teaches the following limitations:

29. The method as recited in claim 24 further comprising, in response to a command indicating an end of the first or second phase, returning to the simulator program.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm"

28. In regards to Claim 33, Agrawal_1 teaches the following limitations:

33. A distributed simulation system comprising a plurality of nodes wherein each node of the plurality of nodes is configured to simulate a different portion of a system under test using a simulator program configured to perform a simulation as a series of timesteps,

and wherein the plurality of nodes are configured to communicate using commands, and a first node of the plurality of nodes is configured to cause the simulator program to evaluate the model in response to receiving a first command including one or more signal values for signals of the model during a first timestep, and

wherein the first node is configured to cause the simulator program to re-evaluate the model in response to receiving a second command including one or more signal values for signals of the model during the first timestep.

See Agrawal_1, p.20, "2.1. The Basic Event-Driven Simulation Algorithm", which teaches a *fan-out update* phase and an *evolution* phase.

Section 2.1 of Agrawal_1 also teaches that "Processing events in two phases permits all input changes at one gate to be accumulated before the gate is evaluated", and

Section 2.1 of Agrawal_1 also teaches that "The basic philosophy of this technique is to *evaluate* (compute the output as a function of the inputs) a gate or functional unit if and only if an *event* (change in the value of an input signal) occurs at the gate inputs."

Claim Rejections - 35 USC § 103

29. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all

obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

30. This application currently names joint inventors. In considering patentability of the claims under 35 U.S.C. 103(a), the examiner presumes that the subject matter of the various claims was commonly owned at the time any inventions covered therein were made absent any evidence to the contrary. Applicant is advised of the obligation under 37 CFR 1.56 to point out the inventor and invention dates of each claim that was not commonly owned at the time a later invention was made in order for the examiner to consider the applicability of 35 U.S.C. 103(c) and potential 35 U.S.C. 102(e), (f) or (g) prior art under 35 U.S.C. 103(a).

31. The prior art used for these rejections is as follows:

- a. Agrawal, P. et al. "A Hardware Logic Simulation System." IEEE Transactions on CAD of Integrated Circuits and Systems. Jan. 1990. Vol. 9, Issue 1, pp.19-29. (Henceforth referred to as "**Agrawal_1**").
- b. Agrawal, P. et al. "Architecture and Design of the MARS Hardware Accelerator." Proc. of the 24th ACM/IEEE Conf. on Design Automation. 1987. pp.101-107. (Henceforth referred to as "**Agrawal_2**").

32. The claim rejections are hereby summarized for Applicant's convenience. The detailed rejections follow.

33. Claims 5-14, 21, 23, 30, and 32 are rejected under 35 U.S.C. 103(a) as being unpatentable over Agrawal_1 in view of Agrawal_2.

In regards to claims 5-14 and 30,

34. In regards to Claim 5,

5. The distributed simulation system as recited in claim 1 further comprising a hub coupled to the plurality of nodes, wherein the hub is configured to receive at least one command from each node during the first phase prior to transmitting commands to the plurality of nodes during the first phase.

Agrawal_1 teaches sending the "zero message" at the end of the evaluation phase "down the pipeline to reconfigure the pipeline for the next fan-out update phase". (See Agrawal_1, p.21, "2.2.2 Evaluation Phase", 1st para.). Examiner interprets that Agrawal_1's "zero message" corresponds to the claimed "no-operation command."

Agrawal_1 also teaches the use of a "crossbar switch" in "a multiple cluster MARS system consist[ing] of several clusters communicating with each

other through a high-speed interconnection network.” (See Agrawal_1, p.19, “I. Introduction”, right column, 3rd para.).

However, Agrawal_1 does not expressly teach that the the “high-speed interconnection network” includes a hub, as recited in claims 5-14 and 30. Agrawal_2, on the other hand, teaches that the crossbar switch in Agrawal_1’s “MARS” system acts as a network hub (see Agrawal_2, Figs.2 and 3, and associated text).

It would have been obvious to modify the teachings of Agrawal_1 with those of Agrawal_2, because they teach the structure and behavior of the same “MARS” system. (See Agrawal_1, p.19, right column, 3rd paragraph, which expressly references Agrawal_2).

35. In regards to Claim 6,

6. The distributed simulation system as recited in claim 5 wherein each node of the plurality of nodes is configured to transmit a no-operation command to the hub if that node has no other command to transmit.

Claim 6 is rejected on the same basis as claim 5.

36. In regards to Claim 7,

7. The distributed simulation system as recited in claim 5 wherein the hub is configured to transmit at least one command to each node of the plurality of nodes.

Claim 7 is rejected on the same basis as claim 5.

37. In regards to Claim 8,

8. The distributed simulation system as recited in claim 7 wherein a first command transmitted by the hub to a first node of the plurality of nodes corresponds to a second command received from one of the plurality of nodes if the second command is routed to the first node; and wherein the first command is a no-operation command otherwise.

Claim 8 is rejected on the same basis as claim 5.

38. In regards to Claim 9,

9. The distributed simulation system as recited in claim 1 further comprising a hub coupled to the plurality of nodes, wherein the hub is configured to receive at least one command from each node during the second phase prior to transmitting commands to the plurality of nodes during the second phase.

Claim 9 is rejected on the same basis as claim 5.

39. In regards to Claim 10,

10. The distributed simulation system as recited in claim 9 wherein each node of the plurality of nodes is configured to transmit a no-operation command to the hub if that node has no other command to transmit.

Claim 10 is rejected on the same basis as claim 5.

40. In regards to Claim 11,

11. The distributed simulation system as recited in claim 9 wherein the hub is configured to transmit at least one command to each node of the plurality of nodes.

Claim 11 is rejected on the same basis as claim 5.

41. In regards to Claim 12,

12. The distributed simulation system as recited in claim 11 wherein a first command transmitted by the hub to a first node of the plurality of nodes corresponds to a second command received from one of the plurality of nodes if the second command is routed to the first node, and wherein the first command is a no-operation command otherwise.

Claim 12 is rejected on the same basis as claim 5.

42. In regards to Claim 13,

13. The distributed simulation system as recited in claim 1 further comprising a hub coupled to the plurality of nodes and configured to signal an end of each of the first phase and the second phase.

Claim 13 is rejected on the same basis as claim 5.

43. In regards to Claim 14,

14. The distributed simulation system as recited in claim 13 wherein the hub is configured to receive at least one command from each node prior to transmitting commands to the plurality of nodes, and wherein the hub is configured to signal an end to one of the first phase or the second phase responsive to receiving a no-operation command from each of the plurality of nodes.

Claim 14 is rejected on the same basis as claim 5.

44. In regards to Claim 21, Agrawal_1 teaches the following limitations:

21. A computer readable medium storing instructions which, when executed on a computer, are configured to signal an end of either a first phase or a second phase of a timestep in a distributed simulation system by transmitting a predefined command indicating an end of the first phase or the second phase to each of a plurality of nodes in the distributed simulation system,

and wherein the instructions are configured to signal the end of either the first phase or the second phase responsive to receiving a nooperation packet from each of the plurality of nodes subsequent to transmitting a command other than a no-operation packet to at least one of the plurality of nodes.

Agrawal_1 teaches sending the “zero message” at the end of the evaluation phase “down the pipeline to reconfigure the pipeline for the next fan-out update phase”. (See Agrawal_1, p.21, “2.2.2 Evaluation Phase”, 1st para.). Examiner interprets that Agrawal_1’s “zero message” corresponds to the claimed “no-operation command.”

Agrawal_1 also teaches the use of a “crossbar switch” in “a multiple cluster MARS system consist[ing] of several clusters communicating with each other through a high-speed interconnection network.” (See Agrawal_1, p.19, “1. Introduction”, right column, 3rd para.).

However, Agrawal_1 does not expressly teach that the “zero message” is specifically the size of a single “packet”, as recited in claims 21, 23, 30, and 32. Agrawal_2, on the other hand, does expressly teach that the size of commands in Agrawal_1’s “MARS” system corresponds to a packet (see Agrawal_2, p.102, left column).

It would have been obvious to modify the teachings of Agrawal_1 with those of Agrawal_2, because they teach the structure and behavior of the same “MARS” system. (See Agrawal_1, p.19, right column, 3rd paragraph, which expressly references Agrawal_2).

45. In regards to Claim 23, Agrawal_1 teaches the following limitations:

23. The computer readable medium as recited in claim 21 wherein the instructions route commands from one of the plurality of nodes to others of the plurality of nodes.

Claim 23 is rejected on the same basis as claim 21.

46. In regards to Claim 30,

30. A method comprising;

signaling an end of a first phase of a timestep in a distributed simulation system by a hub of the distributed simulation system,

the signaling including transmitting a predefined command to each of a plurality of nodes in the distributed simulation system,

wherein signaling the end of the first phase is responsive to receiving a no-operation packet from each of the plurality of nodes subsequent to transmitting a command other than a no-operation packet to at least one of the plurality of nodes.

signaling an end of a second phase of a timestep in a distributed simulation system by the hub,

the signaling including transmitting a predefined command to each of the plurality of nodes in the distributed simulation system.

Agrawal_1 teaches sending the “zero message” at the end of the evaluation phase “down the pipeline to reconfigure the pipeline for the next fan-out update phase”. (See Agrawal_1, p.21, “2.2.2 Evaluation Phase”, 1st para.). Examiner interprets that Agrawal_1’s “zero message” corresponds to the claimed “no-operation command.”

Agrawal_1 also teaches the use of a “crossbar switch” in “a multiple cluster MARS system consist[ing] of several clusters communicating with each other through a high-speed interconnection network.” (See Agrawal_1, p.19, “I. Introduction”, right column, 3rd para.).

However, Agrawal_1 does not expressly teach that the the “high-speed interconnection network” includes a hub, as recited in claims 5-14 and 30.

Agrawal_2, on the other hand, teaches that the crossbar switch in Agrawal_1's "MARS" system acts as a network hub (see Agrawal_2, Figs.2 and 3, and associated text).

Moreover, Agrawal_1 does not expressly teach that the "zero message" is specifically the size of a single "packet", as recited in claims 21, 23, 30, and 32. Agrawal_2, on the other hand, does expressly teach that the size of commands in Agrawal_1's "MARS" system corresponds to a packet (see Agrawal_2, p.102, left column).

It would have been obvious to modify the teachings of Agrawal_1 with those of Agrawal_2, because they teach the structure and behavior of the same "MARS" system. (See Agrawal_1, p.19, right column, 3rd paragraph, which expressly references Agrawal_2).

47. In regards to Claim 32, Agrawal_1 teaches the following limitations:

32. The method as recited in claim 30 wherein signaling the end of the second phase is responsive to receiving a no-operation packet from each of the plurality of nodes subsequent to transmitting a command other than a no-operation packet to at least one of the plurality of nodes.

Claim 32 is rejected on the same basis as claim 30.

Correspondence Information

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Ayal I. Sharon whose telephone number is (571) 272-3714. The examiner can normally be reached on Monday through Thursday, and the first Friday of a biweek, 8:30 am – 5:30 pm.

Art Unit: 2123

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Leo Picard can be reached at (571) 272-3749.

Any response to this office action should be faxed to (571) 273- 8300, or mailed to:

USPTO
P.O. Box 1450
Alexandria, VA 22313-1450

or hand carried to:

USPTO
Customer Service Window
Randolph Building
401 Dulany Street
Alexandria, VA 22314

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to the Tech Center 2100 Receptionist, whose telephone number is (571) 272-2100.

Ayal I. Sharon
Art Unit 2123
December 22, 2005


Paul L. Rodriguez 12/23/05
Primary Examiner
Art Unit 2125